Material Recognition CNNs and Hierarchical Planning for Biped Robot Locomotion on Slippery Terrain

Martim Brandão, Yukitoshi Minami Shiguematsu, Kenji Hashimoto and Atsuo Takanishi

Abstract-In this paper we tackle the problem of visually predicting surface friction for environments with diverse surfaces, and integrating this knowledge into biped robot locomotion planning. The problem is essential for autonomous robot locomotion since diverse surfaces with varying friction abound in the real world, from wood to ceramic tiles, grass or ice, which may cause difficulties or huge energy costs for robot locomotion if not considered. We propose to estimate friction and its uncertainty from visual estimation of material classes using convolutional neural networks, together with probability distribution functions of friction associated with each material. We then robustly integrate the friction predictions into a hierarchical (footstep and full-body) planning method using chance constraints, and optimize the same trajectory costs at both levels of the planning method for consistency. Our solution achieves fully autonomous perception and locomotion on slippery terrain, which considers not only friction and its uncertainty, but also collision, stability and trajectory cost. We show promising friction prediction results in real pictures of outdoor scenarios, and planning experiments on a real robot facing surfaces with different friction.

I. INTRODUCTION

Legged and humanoid robot locomotion planning is an important problem for disaster response and service robots. One of the difficulties of this problem is the complexity of general environments and the need to consider several factors such as collision, energy consumption, surface geometry and friction. In this paper we deal with the specific problem of humanoid robot locomotion when environment friction is considered. Our claim is that friction and its uncertainty can be estimated from vision and robustly integrated into algorithms for motion planning with contact. We argue that even if the precise coefficient of friction cannot be predicted from vision before touching a surface, priors and accumulated experience associated with surface material or condition (think coefficient of friction tables) can provide a probability distribution of friction. Motion planning with contact can also become prohibitively expensive once multiple factors are considered, such as locomotion cost, collision and friction. In this paper we propose a hierarchical approach to the problem, where a footstep planner optimizes the same cost function as a full-body motion planner by use of an oracle, and considers

collision and friction by using simple bounding box collision checks and an "extended footstep planning" [1] approach.

The contributions of this paper are the following:

- We propose a solution to the friction from vision problem using a state-of-the-art deep Convolutional Neural Network (CNN) architecture to predict broad material classes from images, together with known (or learned) distributions of material friction;
- We propose a hierarchical planning architecture for biped robots that optimizes the same objective at both levels, and deals with friction, stability, collision and cost to produce full-body trajectories;
- We show empirical friction prediction results, as well as planning experiments which show the usefulness and applicability of the approach in complex environments with varying friction.

II. RELATED WORK

Recent full-body motion planning [2], [3] and control [2], [4], [5] algorithms for legged robots have started to consider friction by using friction cones in optimization problems. Such methods rely on the fundamental assumption that the coefficient of friction can be predicted in advance. While in itself a challenging problem, partial evidence from human visual perception motivate such an approach to the problem. For example, humans are known to use visual cues to estimate friction, related to surface texture [6], shine [7] and detection of materials or contaminants (e.g. water) [8]. Furthermore, in the human gait literature there is evidence that humans use accumulated previous experience to predict friction and adapt walking style before touching slippery ground [9]. In this paper, such "accumulated previous experience" is implemented as probability distributions of friction associated with material classes. The term "material" is used in a broad sense to refer to visually classifiable classes related to material, condition and context (e.g. "dry metal", "wet asphalt road"). The use of material classes for prediction here is motivated by a recent study [10] which identifies material as one of the most predictive features of both coefficient of friction values and human judgements of friction.

Friction estimation work related to this paper includes that of Angelova et. al [11], which predicts the percentage of slip (i.e. lack of locomotion progress) of a rover from terrain classification and slope. Predictions are made based on non-linear regression of data gathered on a learning stage. Compared to [11], we estimate the coefficients of friction instead of slip, we use open segmentation datasets [12], [13] to train material classification, and we decouple the problem

^{*}This work was supported by JSPS KAKENHI Grant Number 15J06497 and ImPACT TRC Program of Council for Science, Technology and Innovation (Cabinet Office, Government of Japan).

M. Brandão and Y. M. Shiguematsu are with the Graduate School of Advanced Science and Engineering, Waseda University.

K. Hashimoto is with the Waseda Institute for Advanced Study, and is a researcher at the Humanoid Robotics Institute (HRI).

A. Takanishi is with the Department of Modern Mechanical Engineering, Waseda University; and the director of the Humanoid Robotics Institute (HRI), Waseda University.

from the physical robot. Importantly, our approach allows for sharing material friction data among different robots as long as they have similar foot soles.

Several approaches exist to the friction-constrained motion planning problem for legged robots. One approach is the nonhierarchical, full-scale trajectory optimization formulation with implicit contact constraints of [14], [15]. While technically elegant and showing promising results, these can still be computationally expensive for online planning. In order to make the problem tractable, full-body motion can be planned after contact (or footstep) planning [16], [17], in what is called the contact before motion approach. One common issue with such methods is that contact planners do not take the same friction or trajectory criteria into account as the subsequent full-body planners. One exception to this lack of consistency between planning levels is the extended footstep planning work of [1], in which learned models are used at the footstep planning level that predict full-body feasibility. The approach can also account for friction constraints by using timing variables and learned slippage models at the footstep planning level. Still, in [1] the costs optimized at the footstep planning level are not further optimized at the fullbody level. In this paper we improve the method by using an oracle at the footstep planning level which predicts the costs obtained by a full-body trajectory optimizer, thus increasing consistency across planning levels.

III. FRICTION FROM VISION

In this paper we propose to estimate friction of surfaces from visual input by classifying surface material at each image pixel and assuming known (or learned) probability distributions of friction for each material. For convenience we will use the term "friction of a material" to refer to the coefficient of friction between the robot foot sole and a second surface of a given material.

We consider a pixel-wise labelling algorithm that, given an input image I with n pixels, provides a probability distribution $P(X|\theta, I)$, where $X = \{x_1, ..., x_n\}$ are the pixel labels and θ are internal parameters of the algorithm. Each pixel can take one of m possible labels, such that $x_k \in \mathcal{L} = \{l_1, ..., l_m\}$. Furthermore, let each label be a material associated with a probability distribution function (p.d.f.) of a coefficient of friction $p(\mu|l_i)$. Then at pixel k, the conditional p.d.f. of μ is

$$p(\mu|\theta, I) = \sum_{i=1}^{m} p(\mu|l_i) P(x_k = l_i|\theta, I).$$
(1)

For the results shown in this paper we estimated the friction distributions $p(\mu|l_i)$ experimentally, by measuring maximum friction force of the robot foot on several surfaces for each material. We describe the procedure in more detail in Section V.

We use a deep convolutional neural network (CNN) to obtain pixel-wise material predictions $P(x_p = l_i | \theta, I)$. In particular we use the encoder-decoder architecture of [18], which achieves good results in image segmentation applications and is characterized by a low number of parameters. Its



Fig. 1. Our hierarchical planning architecture, which uses trajectory optimization to minimize a cost function $f_{\rm cost}$, as well as oracle costs to plan footstep placement and timing that will have low predicted $f_{\rm cost}$.

low number of parameters leads to fast inference, which is crucial for robotics. The architecture consists of an encoder network of 13 convolutional layers as in VGG16 [19], followed by a decoder network of 13 layers and a final softmax layer. The output of the last layer of the network (a softmax classifier) is at each pixel a vector of probabilities for each class, that is, the probabilities $P(x_p = l_i | \theta, I)$ used in equation (1).

IV. HIERARCHICAL PLANNING

In this paper we plan full-body robot motion using a *contact before motion* approach. A footstep planner first searches a stance graph using transition costs provided by an oracle. The stances are then used as constraints in a full-body trajectory optimizer that considers full-body trajectory costs, collisions, joint limits and static stability. The obtained trajectory is finally interpolated and locally adapted for dynamic stability using a ZMP-based method. The oracle basically takes each stance transition and predicts the costs obtained at the end of the whole planning pipeline. This leads to footstep plans which optimize the same criteria as the full-body planner. See Figure 1 for a visual representation of the architecture.

A. Extended footstep planning with an oracle

The footstep planner searches a graph of stances to find a feasible path between the start and goal stance. Each node in the graph is a stance s, which is defined by a set of contacts with the environment. A contact is a tuple (link, position, rotation). A neighbor stance s' either adds or removes a contact with respect to s. In this paper we deal with biped walking only, and hence stances simply transition from double-support to left-foot-contact, to doublesupport, to right-foot-contact, back to double-support, etc. The advantage of this representation instead of, for example, double-support stances only, is that the swept-volume between consecutive stances can be used by the optimizer to guide a swing leg out of collision. Such an approach is also used by other works focusing on collision detection [20].

In this paper we use the *extended footstep planning* framework of [1]. The footstep planner is "extended" because extra parameters associated with stance transitions (e.g. step timing) are computed from the transition itself by a function

learned offline. Here we call this function an oracle because it predicts the costs that will be obtained by a subsequent full-body trajectory optimizer.

We now briefly describe the footstep planning algorithm. We first discretize the search space by constraining contact positions to a point cloud, and rotations by aligning contact normals with the environment and constraining the links' yaw orientation to a discrete set of values in the global coordinate frame. Then we use an A* variant, ARA* [21], to search the stance graph based on oracle costs. At each state (i.e. stance) of the graph s, a contact is either added or removed to generate successor stances. Contact removal generates one new stance. On the other hand, adding a new contact consists of doing a range search of points in a radius around the foot in contact. For each of those points, footsteps are placed at all yaw angles and checked for feasibility. We implement feasibility as empirical stance distance limits, as well as foot-foot and COM-environment collision checking using bounding boxes for the feet and trunk. The feasible stances are added as successors of s.

To find an optimal path to the goal state, A* search requires a state transition cost function c(s, s') and a heuristic cost-to-go function h(s). Similarly to [1], we define the cost as

$$c(s, s') = \min_{p} \hat{f}_{cost}(s, s', p)$$

subject to
$$P(\hat{f}_{RCOF}(s, s', p) < \mu^{(k)}) \ge \eta$$

$$k = 1, ..., K,$$
(2)

where k is an index of the contacts of s and s', $\mu^{(k)}$ is the friction at these contacts, and p are state transition parameters. Since in this paper we use two full-body posture waypoints per stance at the trajectory optimization level (Section IV), we set transition parameters $p = (\Delta t, \Delta t')$. These are the time spent from the second wavpoint of s until the first waypoint of s', and the time spent from the first waypoint of s' to the second, respectively. The main difference in (2) with respect to [1] is that we consider uncertainty in the coefficient of friction variable by using chance constraints. RCOF stands for required coefficient of friction and corresponds to the maximum tangential-tonormal force ratio exerted over the whole trajectory [1]. Therefore, the constraints $P(\hat{f}_{\text{RCOF}}(s, s', p) < \mu^{(k)}) \geq \eta$ in (2) implement robust Coulomb friction at each contact by forcing the inequalities to hold with at least probability η . The constraints can also be rewritten using the cumulative distribution function of (1) denoted by $F_{\mu^{(k)}|\theta,I}$,

$$F_{\mu^{(k)}|\theta,I}(\hat{f}_{\text{RCOF}}(s,s',p)) \le 1 - \eta.$$
 (3)

Since each $\mu^{(k)}$ is one-dimensional then F can be inverted and the constraints rewritten in deterministic form

$$\hat{f}_{\text{RCOF}}(s, s', p) \le Q_{1-\eta}^{(k)},$$
 (4)

where $Q_{1-\eta}^{(k)}$ is the $(1-\eta)$ -quantile of $F_{\mu^{(k)}|\theta,I}$, which can be computed by an integral of (1) over μ .

Regarding the heuristic cost-to-go function of A^* search, as in [1], we set it to

$$h(s) = d_{xy}(s, s^{\text{goal}}) \cdot \min_{s, s', p} \frac{\hat{f}_{\text{cost}}(s, s', p)}{d_{xy}(s, s')},$$
(5)

where the function $d_{xy}(.,.)$ computes Euclidean distance on the horizontal plane between two stances (i.e. the distance between left feet and right feet summed). The heuristic (5) is a lower bound on the cost-of-transport times distance, which guarantees that h(s) does not overestimate the total cost to the final stance s^{goal} (i.e. is admissible, a necessary condition for A* optimality).

In this paper, the functions f_{cost} and f_{RCOF} are given by an oracle which predicts the value of f_{cost} and f_{RCOF} obtained at the end of the whole planning pipeline. Notice that in this paper, contrary to [1], f_{cost} is the function that will be optimized at the full-body trajectory optimization level. We implement \hat{f}_{cost} and \hat{f}_{RCOF} as hash tables. The tables are filled offline, by feeding the whole planning pipeline (i.e. trajectory optimization, interpolation, dynamic stabilization) with uniformly distributed samples of (s, s', p)as shown in Figure 1. The discrete optimization problems in (2), (5) are then solved for a large number of discretized stances and coefficient of friction quantiles and stored in new hash tables for fast access to costs and heuristics during search.

B. Full-body trajectory optimization

The full-body trajectory optimizer takes a footstep plan with N stances and produces a full-body trajectory, parameterized by T discrete-time waypoints. Waypoints are fullbody robot configurations $q_t \in \mathbb{R}^D$, t = 1, ..., T, where D is the number of degrees-of-freedom consisting of the joints' angle values and the pose of the robot base. Each stance is associated with 2 full-body postures (at start and midstance) and so T = 2N. For convenience we use s_t to refer to the stance associated to q_t .

Our optimizer solves the problem

 $\underset{q_1,...,q_T}{\text{minimize}} \quad f_{\text{cost}}(q_1,...,q_T) + \alpha f_{\text{collision}}(q_1,...,q_T) \quad (6a)$

subject to

$$f_{\text{stance}}(q_t, s_t) = 0 \quad \forall_{t \in 1, \dots, T}$$
(6b)

$$f_{xy}(q_t) \in \mathcal{P}_t \quad \forall_{t \in 1, \dots, T} \tag{6c}$$

$$f_{\text{roll}}(q_t) = 0 \quad \forall_{t \in 1, \dots, T} \tag{6d}$$

$$A_t q_t \le b_t \quad \forall_{t \in 1, \dots, T},\tag{6e}$$

where $q_1, ..., q_T$ are the optimization variables, α is a penalty constant and:

- The function f_{cost} computes the sum of the squared static torques of all joints at all waypoints, as implemented in the *trajopt* library [22]
- The function $f_{\text{collision}}$ is a collision cost as proposed by [22]. It is the sum of a discrete collision cost computed by the signed distance between each link and all other geometries, and a continuous collision cost computed

by the signed distance between the swept volume of each link with the environment

- The function $f_{\text{stance}}(q_t, s_t)$ computes the pose error of all links in contact as a 6*C*-dimensional vector where C is the number of active contacts in s_t . This is computed as the translation and axis-angle error between the target link pose (given by s_t) and the current link pose (given by q_t)
- The function $f_{xy}(q_t)$ computes the (x,y) coordinates of the COM, and \mathcal{P}_t is the support polygon of s_t . The constraint thus enforces approximate static stability. The support polygon of s_t is computed by the convex hull of the horizontal projection of links in contact and does not include contacts removed in s_{t+1}
- The function $f_{\rm roll}(q_t)$ computes the rotation around the X axis for the waist link, with respect to the global reference frame. This constraint is necessary as "zero roll" is an assumption of the subsequent dynamic stabilization method (Section IV-C)
- A_t, b_t enforce joint angle and velocity limits.

We solve problem (6) using the Sequential Quadratic Programming method of [22] as implemented in the *trajopt* library¹.

C. Interpolation and stabilization

To obtain a densely-sampled trajectory for execution on the robot, we interpolate trajectory waypoints using Hermite cubic splines with derivatives set to zero for smooth contact transitions. The time between two consecutive waypoints q_t is given by the oracle, as we describe in Section IV-A.

Since the obtained trajectory is not dynamically stable, we then apply an FFT-based ZMP trajectory compensation scheme [23]. The method considers the rigid-body dynamics of the full body and locally adapts COM motion on the horizontal plane using analytic inverse kinematics to iteratively reduce the error between the real and reference ZMP trajectory. We set the reference ZMP trajectory to the interpolated $f_{xy}(q_t)$, which were used in the optimization problem (6) and are inside the support polygon at each waypoint. Furthermore, our implementation of the analytic inverse kinematics of the robot WABIAN-2 assumes zero roll angle of the waist link with respect to the world reference frame. We include this constraint in the optimization problem (6) for consistency.

V. RESULTS

A. Material segmentation results

To train the CNN we first collected 7,791 annotated images from publicly available semantic-segmentation datasets: 5,216 from the VOC2010 Context dataset [13] and 2,575 from the OpenSurfaces dataset [12]. We selected all images in the datasets with at least one of the following labels: *asphalt, concrete, road, grass, rock, sand, sky, snow, water, carpet, rug, mat, ceramic, tile, cloth, fabric, marble, metal, paper, tissue, cardboard, wood.* Due to similarity between some classes at the image and semantic level we joined the labels (*asphalt, concrete, road*), (*carpet, rug, mat*), (*ceramic, tile*), (*cloth, fabric*) and (*paper, tissue, cardboard*). The total number of considered classes in the output CNN layer was 14. *Sky* was only included to avoid classifying it as any of the other materials on outdoor pictures.

We used stochastic gradient descent with 0.1 learning rate and 0.9 momentum as in the original SegNet publication [18], and trained the network on an Amazon Elastic Cloud node with a 4GB NVIDIA GPU. We ran a total of 90,000 iterations with a mini-batch size of 5 (maximum allowed by the GPU). Training was done on 60% of the images, while the other 40% were used as the test set.

We obtained a global classification accuracy of 0.7929 and class-average accuracy of 0.4776 on the test set. See Figure 2 for examples of the (highest probability) material predictions given by the CNN on the test set. The global accuracy is comparable to state-of-the-art performance in semantic segmentation (e.g. [18], [24]), and the class-average accuracy is slightly below state-of-the-art (which is around 0.60 [18]). We believe one important way to improve classification accuracy is to improve the dataset itself since, for instance, there is moderate visual similarity between some of the materials such as marble and ceramic, and some materials are lowly sampled (e.g. the lowest sampled materials are snow and sand, present in 173 and 46 images respectively).

The material segmentation results in Figure 2 show an overall good accuracy of the CNN, particularly on wood, grass and sky labels. The figure also shows typical misclassifications such as white walls recognized as sky or metal (picture 6), hard snow as rock (picture 7), and some overlap between asphalt/road, ceramic and marble. These are arguably understandable since material labels themselves semantically overlap. However, our approach to the visual friction estimation problem is such that if there is uncertainty in the material label, then this uncertainty can be used to weight the friction of the surface through material and friction probability distributions (Section III).

B. Friction prediction results

We empirically measured the coefficient of friction associated with each material label using a force gauge and the robot foot loaded with a 1.5kg mass. The foot is rigid and its sole is covered with a high stiffness soft material for shock absorption and an anti-slippage sheet. We checked whether surfaces were horizontal with a level, then placed the foot and measured maximum friction force with the force gauge. See Figure 3 for an illustration of the procedure. We took 5 friction measurements on each surface, and used at least 3 surfaces of each material. We fitted a normal distribution to the measurements, obtaining separate parameters μ_i and σ_i^2 for each material, where μ_i is the mean friction of material *i*, and σ_i^2 the variance. The materials sand, snow, water, cloth, paper were an exception, and since our robot is currently not capable of walking on them (i.e. fall or damage risk is too high) we directly set them to $\mu_i = 0, \sigma_i^2 = 0$. We

¹URL: http://rll.berkeley.edu/trajopt



Asphalt Grass Rock Sand Sky Snow Water Carpet Ceramic Cloth Marble Metal Paper Wood

Fig. 2. Example images from the test set, respective material segmentation (highest probability labels) and coefficient of friction quantiles $Q_{1-0.95}$ of equation (1). Darker shades of gray correspond to higher friction, such that white is $\mu = 0$ and black $\mu = 1$.



Fig. 3. We estimate the coefficient of friction of a material by several measurements of the maximum friction force on the robot foot, loaded with a 1.5kg mass.

similarly set *sky*'s friction to zero as well. See Table I for the parameters of the friction p.d.f. of each material.

In Figure 2 we show the test-set's highest probability material predictions along with the $(1 - \eta)$ -quantile of the

 TABLE I

 NORMAL DISTRIBUTION PARAMETERS OF EACH MATERIAL'S

 COEFFICIENT OF FRICTION, MEASURED MANUALLY ON THE ROBOT FOOT

Material	μ_i	σ_i
Asphalt	0.74	0.12
Grass	0.53	0.10
Rock	0.80	0.08
Carpet	0.82	0.02
Ceramic	0.97	0.05
Marble	0.83	0.15
Metal	0.80	0.15
Wood	0.88	0.12
Sand, Sky, Snow, Water, Cloth, Paper*	0	0

Note: materials marked with a * are assumed to be untraversable by our robot and so set to zero without measurements.



Asphalt Grass Rock Sand Sky Snow Water Carpet Ceramic Cloth Marble Metal Paper Wood

Fig. 4. Left: the view from the robot's camera of the mock-up scenario built in the laboratory. The locomotion target is one meter ahead, marked with a white "x". Middle: material segmentation. Right: coefficient of friction quantiles $Q_{1-0.95}$ of equation (1). Darker shades of gray correspond to higher friction, such that white is $\mu = 0$ and black $\mu = 1$.

coefficient of friction which is used in equation (4). We set a typical value of $\eta = 0.95$. The friction images are darker where friction is higher ($\mu = 1$ would be black). Note that ceramic-like surfaces have high predicted friction (pictures 1, 5, 6); beds and jackets have very low friction (pictures 3, 4, 14); grass patches have lower friction than roads (pictures 2, 9, 11, 12, 13); and that water is mostly white - zero friction - (pictures 10, 12).

The figure also shows the advantage of using the whole probability distribution of materials (instead of using the highest probability material) to estimate friction. For example in picture 14, the jacket on the ground is classified as cloth and rock depending on the region, but friction is low on most of the object's area since the *cloth* label still has high probability.

C. Planning results

We prepared a mock-up scenario in the laboratory which demonstrates the capabilities of our planner. The scenario consists of a floor with two areas of different materials. One is made of wood ($\mu = 0.84$) and the other is a high-friction flooring resembling ceramic tiles both in appearance and coefficient of friction ($\mu = 1.00$). The perception-planning algorithms were run on this scenario, and then a piece of cloth (T-shirt) was laid flat on one of the surfaces to provoke changes in friction and force a different plan. See Figure 4 for the scenario, segmentation and friction as seen from the robot's camera at the initial condition. Once again, the figures show the advantage of using the full probability distribution of materials given by the CNN. While *cloth* is the highestranking material only in part of the object region, friction is low on a larger region which is highly consistent with object borders.

The robot starts in double-support, with one foot on each surface. The goal stance is one meter ahead, also with a foot on each surface. After the robot is placed at the initial state, the perception and planning algorithms run without any human input except the push of a button to execute the planned full-body trajectory open-loop. Trajectory optimization parameters are the collision penalty weight α of equation (6), which is set to 50, and the distance at which the collision penalty starts being applied (for all links except those in contact), which we set to 2.5cm. The obtained fullbody trajectory is tracked by position control at the joint level.

For these experiments we used the human-sized humanoid robot WABIAN-2 [25] customized with a Carnegie Robotics Multisense SL sensor-head. The perception pipeline predicts pixel-wise material label distributions and pixel-wise friction using SegNet [18] and equation (1), and combines them with the stereo depth maps computed onboard by the Multisense. It produces friction-annotated point clouds at 2Hz. For collision checking, the point cloud is converted into a mesh using the fast surface reconstruction algorithm of [26] as implemented in PCL [27]. All perception and planning computation ran on an external PC with network connection to the robot's onboard PC, and we used ROS [28] for communication.

We show the results of the perception-planning experiments in Figure 5. From left to right, we show the material and friction point clouds, the footstep plan, the collisionchecking bounding boxes used by the footstep planner and the final planned full-body trajectory after optimization and stabilization. In the first situation there are only wood and ceramic surfaces, but the predicted lower bound of friction of the wood surface is lower than that of the tiles $(Q_{1-0.95} = 0.1)$ vs 0.4). The footstep planner returns a sequence of stances that reduces the amount of times wood is stepped on. This behavior comes naturally from the extended footstep planning approach [1], since walking on low friction ground requires higher stance times (slower motion) and thus more energy cost. Furthermore, note that the trajectory optimization uses all degrees-of-freedom to satisfy the constraints (e.g. trunk roll use is clear in the image sequence, important mainly for the stability constraints), and that the knees are relatively stretched in order to reduce torque consumption but still satisfy stability constraints. Also note that swing leg clearance happens automatically due to the use of collision costs.

In the second situation we laid a flat piece of cloth on a ceramic spot used by the previous trajectory. The cloth was correctly classified and its friction was practically zero. The footstep planner returned a trajectory around the cloth and on the wood surface, which led to a slightly longer time and energy cost of the full-body trajectory (63 vs 60 seconds, 5% longer than on the first situation). Note that while the times are long they correspond to 25 stances because of step length limits, and thus the average time per stance is approximately 2.5 seconds.

Footstep planning took approximately 20 seconds in the first situation and 10 in the second. The reason for the difference is clear from the scenario: while in the first situation stances on both surfaces are expanded by A* in order to guarantee optimality, in the second situation no stances are expanded on the surface with cloth since friction zero has infinite cost. Full-body trajectory optimization



Fig. 5. Perception and planning experiments with two surfaces (wood and ceramic) on the first two rows and three surfaces (wood, ceramic and cloth) on the last two rows. We show the material segmentation (same colour codes as in Figure 4), friction (cold colours are high friction, warm are low), footstep plan, collision bounding boxes, full-body plan and finally the walking sequence on the real robot.

took approximately 40 seconds and dynamic stabilization 2 seconds. Note that these are for 25-stance, 60 second trajectories, and therefore they should be considerably faster in case planning is done one or two steps at a time.

VI. CONCLUSIONS AND DISCUSSION

In this paper we proposed a complete solution to the problem of biped robot locomotion on slippery terrain. We developed both a visual friction estimation algorithm and an objective-consistent hierarchical planning method which considers trajectory costs, collision, stability and friction.

We empirically showed that friction estimates in our algorithm are more consistent with object/material borders than the highest-probability material label segmentation, which shows a good integration of segmentation uncertainty into friction estimation. We also showed that the algorithms work for varied terrain and are applicable to planning on a real robot. The algorithms are relevant since not only obstacles but also different terrain types abound in the real world, and locomotion choices should take them into account - whether for safety or energetic considerations.

Regarding the perception problem, we opted to decouple it into (broad sense) material segmentation and per-material friction distributions. Even though we obtained the material friction distributions manually, these could also be learned over time with locomotion experience. Alternatively, friction could also be learned from images directly by end-to-end training, for example by initialization of a CNN with the parameters obtained with our architecture.

For the context of this paper all surfaces were dry. Wet surfaces could also be included, although from our experience they should be treated as separate material labels (e.g. "dry metal" and "wet metal") so that the distribution $\mu | l_i$ does not become bimodal. Thus, one important detail in this work is the notion of material, which should be taken in a broad sense, as a visually distinguishable terrain class.

Importantly, one problem with the proposed perception and planning approach is that wrong material classifications can lead to there being no solution to the footstep planning problem. An example of such a situation is when a material the robot cannot walk on, such as water in our case, is mistakenly given very high confidence. Our view is that the solution could be semi-supervision where a teleoperator can correct a segmented region's material label. However, we believe that humans should not directly annotate COF, since despite their relative ability to adapt gait to slippery ground humans have difficulties in estimating coefficient of friction values [6], [10].

Finally, full-body trajectories in this paper were interpolated after trajectory optimization at waypoints. While we did this for implementation simplicity, one possible direction of improvement could be to use the spline representation directly in the optimization problem, using constraints at collocation points.

REFERENCES

- M. Brandao, K. Hashimoto, J. Santos-Victor, and A. Takanishi, "Footstep planning for slippery and slanted terrain using human-inspired models," *IEEE Transactions on Robotics*, vol. 32, no. 4, pp. 868–879, Aug 2016.
- [2] S. Kuindersma, R. Deits, M. Fallon, A. Valenzuela, H. Dai, F. Permenter, T. Koolen, P. Marion, and R. Tedrake, "Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot," *Autonomous Robots*, vol. 40, no. 3, pp. 429–455, 2016.

- [3] K. Hauser, T. Bretl, J.-C. Latombe, K. Harada, and B. Wilcox, "Motion planning for legged robots on varied terrain," *The International Journal of Robotics Research*, vol. 27, no. 11-12, pp. 1325–1349, 2008.
- [4] A. Herzog, N. Rotella, S. Mason, F. Grimminger, S. Schaal, and L. Righetti, "Momentum control with hierarchical inverse dynamics on a torque-controlled humanoid," *Autonomous Robots*, vol. 40, no. 3, pp. 473–491, 2016.
- [5] S. Feng, X. Xinjilefu, W. Huang, and C. Atkeson, "3d walking based on online optimization," in 13th IEEE-RAS International Conference on Humanoid Robots, October 2013.
- [6] M. F. Lesch, W.-R. Chang, and C.-C. Chang, "Visually based perceptions of slipperiness: Underlying cues, consistency and relationship to coefficient of friction," *Ergonomics*, vol. 51, no. 12, pp. 1973–1983, 2008, pMID: 19034787.
- [7] A. S. Joh, K. E. Adolph, M. R. Campbell, and M. A. Eppler, "Why walkers slip: Shine is not a reliable cue for slippery ground," *Perception & Psychophysics*, vol. 68, no. 3, pp. 339–352, 2006.
- [8] K. W. Li, W.-R. Chang, T. B. Leamon, and C. J. Chen, "Floor slipperiness measurement: friction coefficient, roughness of floors, and subjective perception under spillage conditions," *Safety Science*, vol. 42, no. 6, pp. 547 – 565, 2004.
- [9] G. Cappellini, Y. P. Ivanenko, N. Dominici, R. E. Poppele, and F. Lacquaniti, "Motor patterns during walking on a slippery walkway," *Journal of Neurophysiology*, vol. 103, no. 2, pp. 746–760, 2010.
- [10] M. Brandao, K. Hashimoto, and A. Takanishi, "Friction from vision: A study of algorithmic and human performance with consequences for robot perception and teleoperation," in 2016 IEEE-RAS International Conference on Humanoid Robots, Nov 2016.
- [11] A. Angelova, L. Matthies, D. Helmick, and P. Perona, "Slip prediction using visual information," in *Proceedings of Robotics: Science and Systems*, Philadelphia, USA, August 2006.
- [12] S. Bell, P. Upchurch, N. Snavely, and K. Bala, "OpenSurfaces: A richly annotated catalog of surface appearance," ACM Trans. on Graphics (SIGGRAPH), vol. 32, no. 4, 2013.
- [13] R. Mottaghi, X. Chen, X. Liu, N.-G. Cho, S.-W. Lee, S. Fidler, R. Urtasun, and A. Yuille, "The role of context for object detection and semantic segmentation in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [14] M. Posa and R. Tedrake, Direct Trajectory Optimization of Rigid Body Dynamical Systems through Contact. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 527–542.
- [15] I. Mordatch, E. Todorov, and Z. Popović, "Discovery of complex behaviors through contact-invariant optimization," ACM Trans. Graph., vol. 31, no. 4, pp. 43:1–43:8, July 2012.

- [16] R. Deits and R. Tedrake, "Footstep planning on uneven terrain with mixed-integer convex optimization," in 14th IEEE-RAS International Conference on Humanoid Robots, Nov 2014, pp. 279–286.
- [17] A. Escande, A. Kheddar, and S. Miossec, "Planning contact points for humanoid robots," *Robotics and Autonomous Systems*, vol. 61, no. 5, pp. 428 – 442, 2013.
- [18] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," arXiv preprint arXiv:1511.00561, 2015.
- [19] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014.
- [20] N. Perrin, O. Stasse, L. Baudouin, F. Lamiraux, and E. Yoshida, "Fast humanoid robot collision-free footstep planning using swept volume approximations," *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 427–439, April 2012.
- [21] M. Likhachev, G. J. Gordon, and S. Thrun, "Ara*: Anytime a* with provable bounds on sub-optimality," in Advances in Neural Information Processing Systems, 2003, pp. 767–774.
- [22] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [23] K. Hashimoto, H. Kondo, H.-O. Lim, and A. Takanishi, Motion and Operation Planning of Robotic Systems: Background and Practical Approaches. Springer International Publishing, 2015, ch. Online Walking Pattern Generation Using FFT for Humanoid Robots, pp. 417–438.
- [24] S. Bell, P. Upchurch, N. Snavely, and K. Bala, "Material recognition in the wild with the materials in context database," *Computer Vision* and Pattern Recognition (CVPR), 2015.
- [25] Y. Ogura, H. Aikawa, K. Shimomura, H. Kondo, A. Morishima, H. Lim, and A. Takanishi, "Development of a new humanoid robot wabian-2," in 2006 IEEE/RSJ International Conference on Robotics and Automation. IEEE-RAS, 2006.
- [26] D. Holz and S. Behnke, Fast Range Image Segmentation and Smoothing Using Approximate Surface Reconstruction and Region Growing. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 61–73.
- [27] R. B. Rusu and S. Cousins, "3D is here: Point Cloud Library (PCL)," in 2011 IEEE International Conference on Robotics and Automation, Shanghai, China, May 9-13 2011.
- [28] M. Quigley, K. Conley, B. P. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.